

ABSTRACT

Read-only and read-write snapshot copies of a production file in a Unix-based file system are organized as a version set of file inodes and shared file blocks. Version pointers and branch pointers link the inodes. Initially the production file can have all its blocks preallocated or it can be a sparse file having only an inode and its last data block. A protocol is provided for creating read-only and read-write snapshots, deleting snapshots, restoring the production file with a specified snapshot, refreshing a specified snapshot, and naming the snapshots. Block pointers are marked with a flag indicating whether or not the pointed-to block is owned by the parent inode. A non-owner marking is inherited by all of the block's descendants. The block ownership controls the copying of indirect blocks when writing to the production file, and also controls deallocation and passing of blocks when deleting a read-only snapshot.